

Towards Adaptive Quest Narrative in Shared, Persistent Virtual Worlds

Emmett Tomai

University of Texas – Pan American, 1201 W. University Dr., Edinburg, TX 78539
tomaie@utpa.edu

Abstract

In this paper, we discuss motivations for studying interactive narrative in shared, persistent worlds using the established conventions of quest-based MMORPGs. We present a framework for categorizing the various techniques used in these games according to the interaction between the *world model* and the *quest model*. Using this framework we generalize recent games to present a more dynamic world model, and investigate extensions to the quest model to support storytelling through adaptive quest narratives.

Introduction

Human interaction in shared, persistent virtual spaces has greatly increased in the past decade. The high profile coverage of Linden Lab's *Second Life* virtual world, and the explosive success of Blizzard Entertainment's Massively Multiplayer Online Role-Playing Game (MMORPG) *World of Warcraft (WoW)* are obvious examples. But Ducheneaut, et al (2006) found that actual joint activities in *WoW* were not all that prevalent. In contrast to traditional multiplayer activities, players seemed drawn to be "alone together". They explained this phenomenon as a social desire for audience, social presence and spectacle. The even more explosive social networking phenomena of Facebook appears to show the same desire to feel surrounded by people, but not obligated to interact with them directly. Ducheneaut, et al and Reeves, et al (2005) both conclude that design of human computer experiences should pay more attention to this apparently powerful desire.

Interactive narrative research is concerned with how to use computational models to provide users with experiences that are interactive and yet have the characteristics of narrative. Little work in this field has tackled the challenge of storytelling in a shared, persistent world. MMORPGs provide functional models of such

storytelling, though with significant limitations on interactivity and narrative.

As in other many other genres, narrative is used in MMORPGs to guide and motivate player actions. Due to the popularity of *WoW*, and Sony Online's *EverQuest* before it, there is a firmly entrenched standard for narrative that many games adopt: the *quest* system. Quests combine in-world objectives for players to pursue with narrative discourse that motivates the action. The stories are largely trite and formulaic, yet they are not "toy" narratives. They are legitimate, human authored artifacts that have been consumed by millions of players. Story arcs are built out of many quests, typically in linear chains. As building blocks of narrative, quests have several interesting properties. First, they are modular, allowing players to interleave story arcs together on-demand. Second, they directly tie short free form narrative discourse to computationally friendly structured representations of entities and actions in the world. Third, there are tens of thousands of them, publically available through data mining websites for a number of different MMORPGs. Working in the context of MMORPG raises different challenges for interactive narrative research. Shared persistence lends a weight to the world: it cannot simply change to accommodate any one player. This de-emphasizes the problem of responding to every action a player can come up with, and emphasizes generating or adapting narrative to fit the world situation. In addition, starting with quest systems is an attempt to keep this work connected to actual, playable experiences.

In standard quest systems, authored narratives are combined with a static, unchanging world. In this paper, we propose a framework for analyzing how existing games design the interactions between the *world model* and the *quest model*. In particular, we look at how recent games have attempted to make a more dynamic world that responds to player actions. We then present preliminary work exploring how extensions to those standard world and quest models can allow human authored narratives to be adapted to fit such a changing world.

Related Work

Extending quest conventions for adapting interactive narrative is not a new idea. Li and Riedl (2010) use plotline adaptation, a planning approach that modifies human authored plotlines made up of quests. These quests are represented as abstract planning operators with causal and temporal information. The modifications are done using plan refinement with decomposition and partial order causal link repair. However, while the aesthetic value of the human authored plotline is pointed out, actual narrative discourse at the quest level is not used. It is unlikely that such text would maintain coherence through the adaptation. The GrailGM system (Sullivan, et al. 2011) dynamically chooses quests to give a player in the single-player RPG *Mismanor*. These quests are annotated with hard and soft pre-conditions for selection, encoding authorial knowledge and preference for the selection process. The focus is on providing multiple paths through the same story. In this work, we also look at quest selection constraints, but the shared world constraint significantly changes the problem. Instead of multiple paths through the same story, we are investigating how a variety of loosely related stories can be adapted to work together in a shared, changing world.

Fairclough & Cunningham (2003) and Riedl, et al. (2011) both use the multiplayer MMORPG quest model, but focus on the constraints of automatic plot generation between a few directly interacting players. The True Story system (Pita, et al. 2007) proposes and implements a quest generation system intended for MMORPGs, but it is not evaluated, and it is not clear how the quests would adapt to change in the world. There are many other planning and case-based interactive story generation systems (cf. Riedl 2008; Gervás et al. 2005), but they work with logical representations that are realized through Natural Language Generation or physical action in a virtual environment (cf. Young et al. 2004). By contrast, we begin with authored text already in place. Winegarden and Young (2006) extended such a planning system to maintain consistency between multiple, unrelated player stories in a shared environment. This may be applicable when we turn to more algorithmic concerns.

A Framework for Quest Systems

In a quest-based MMORPG, the quest model works with the model of the simulated world to guide player experience. In this section, we define those models for the basic, common elements of standard quest systems, and then use that framework to analyze newer quest-based designs. In the standard world model, there are large zones divided up into designer-created regions. Each region contains *camps* populated with entities that the player can

interact with, along with terrain, foliage, buildings and other static props. There are mobile, autonomous entities known as *mobs*, which exist for players to fight against. There are also stationary, non-agentive entities known as *nodes*. Interacting with nodes can grant *items* (e.g. picking herbs from a bush) or trigger authored scripts (e.g. open a door). Items are also *dropped* by entities after they are killed. Mobs and nodes are usually controlled by spawning logic that maintains a max number of them in the camp. They may also be controlled by authored scripts, such as spawning in to ambush a player. Finally, there are *non-player characters (NPCs)* for friendly interactions (e.g. trading), often in safe town-type areas.

The standard quest model works with the world model to guide players to camps for episodic encounters. Players interact with *quest givers* to obtain quests, which have camp-specific objectives to be completed at the player's discretion. The most common objectives are killing mobs, obtaining items, moving to locations and interacting with nodes and NPCs. Once the objectives for a quest are complete, the player can interact with a specified *quest completer* in order to *turn-in* the quest and receive rewards. Quests are often organized in chains where one quest must be turned-in before the next can be obtained. Quest givers and completers are often NPCs, but may also be nodes or items. Each quest also contains narrative discourse to provide background and motivate the objectives, presented upon interaction with the quest giver. Additional discourse is presented on turn-in to comment on the resolution.

The standard world model is necessarily unchanging. The camps must always have the same entities available to players, or the standard quest model would not work. This creates incoherence between the narrative discourse that describes player impact, and the world that never changes. Nevertheless, this limitation has not stopped most high-profile MMORPGs from adopting the standard quest model. However, several recent games that have adopted it have made significant changes in an attempt to increase dynamism and player impact on the world.

EA/Mythic's *Warhammer Online (WAR)* added *public quests*. These involve special camps that cycle through multiple stages. For example, there may be weak enemies to begin with, but if they are defeated, more powerful enemies arrive, and so on. Quests for these camps automatically deliver objectives and short, descriptive text to all players nearby whenever a new stage begins, and they either complete or fail when the stage ends. Trion Worlds' *Rift* goes a step further, supplementing standard quests with *invasion* events. Special invading mobs spawn in an area, potentially disrupting other camps and the quests they support. Objectives and descriptions are pushed to players in the area. Unlike public quests, these invasions do not follow scripted stages and can spread across the map if unchecked. ArenaNet's *Guild Wars 2*

(GW2) replaces standard quests with *heart quests*. All quest objectives are automatically shown to players when they are at the relevant camp, and only when they are there. When the objectives are completed, rewards are also pushed to the player. The camps still have NPCs in a quest-giver role, allowing players the option of talking to them for discourse if they wish. This alternative model allows GW2 to have regions that transition between different states, each with its own camps and sets of quests. Events occur in each state, and the next transition is based on whether players win or lose the events. However, this model supports only a subset of the quest discourse seen in other games. ArenaNet appears to still value in-depth narrative, as GW2 uses traditional quest delivery in a player's *personal story*, essentially an embedded single player experience.

Modeling Adaptive Quest Narrative

Each of the prior examples extends the standard world model to add dynamic, player-driven change. However, to work with those extensions, they all adopt quest models that diminish the role of narrative discourse. In this section, we explore how techniques for adapting quest narratives can allow such a changing world to work with the level of narrative found in the standard quest system. In particular, we are concerned with maintaining *coverage*, *coherence* and *usage*. Coverage means that quests are available to players and able to be completed as expected. Coherence means that the narrative discourse in a quest is coherent with what is happening in the world and with other quests. Usage means that quests are actually seen by players. To perform this exploration, we analyzed roughly 50 quests in *WoW* and *Rift* to see what world states and changes they describe, and how they might be adapted to fit a changing world. We also created a sample zone with roughly thirty quests in four stories to work with. We first detail an extended world model that generalizes the features seen in the prior example games. We then use that model to explore extensions to the quest model to maintain coverage, coherence and usage through adaptation.

Extended World Model

In this generalized model, as in the examples, the world is organized into camps that spawn mobs and nodes, and those camps can transition through multiple stages. We focus on combat-related stages, because that is the dominant facet in these games, but there is no reason to believe simple, game models of economics, politics, relationships, etc could not also apply. The primary stages are *foothold* and *established*. In the foothold stage, the camp has few mobs relative to its specified size. Organized mobs in a foothold cluster together in one area,

while animals are spread out sparsely. In the established stage, mobs are spawned to meet their specified numbers and fill the area. In addition, a region may have two camps in conflict, spawning mobs to fight each other. This is a common setup in recent games because so many quests describe ongoing battles. Using the two primary stages, the world can display two small forces in combat (foothold vs. foothold), a small force holding a small area while fighting a large force (foothold vs. established), or two large forces mixed together in combat (established vs. established). Figure 1 shows the transitions between stages. The primary stages are shown as circles, and the conflict stages are shown as rectangles.

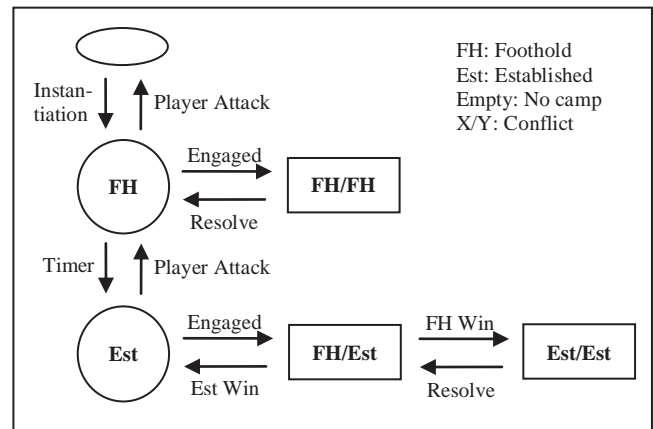


Figure 1. Transitions between camp stages.

Camps are instantiated into the foothold stage (FH) from the empty circle indicating no camp. If a foothold lasts a certain amount of time, it transitions to the established stage (Est). Player attacks can cause transitions back from established to foothold to none. A second camp may be instantiated in an occupied region, engaging an existing camp. The engaging camp is always instantiated as a foothold. If a foothold camp is engaged, this results in two footholds in conflict (FH/FH). This can only be resolved by one overwhelming the other, which returns the region to a single foothold camp. Similar transitions are shown for engaging an established camp. The specific conditions for triggering transitions (e.g. kill a certain number, kill at a certain rate) must be tuned for player experience. Camps in conflict would reach natural equilibrium, but the point is for players to impact the outcome by getting involved.

Mob Types	Mobs that make up the camp
Node Types	Nodes that are spawned in the camp
Props	Non-interactable objects in the area
Size	Qualitative indication of the max spawn level
Location	Region the camp is instantiated in
Stage	Foothold or Established
Engaged	True or false

Table 1. Camp attributes.

Camps also have specific nodes and props in them, as in the standard world model. The complete state of each camp is described by the attributes shown in Table 1. These attributes control spawn behavior and interact with the quest model for adaptation.

Extended Quest Model

In the standard models, the unchanging world has a number of camps each in a fixed state. Every quest is linked to one of those camps, and can be assigned and completed at any time. One way to view the challenge of a changing world model is that the state space is considerably increased. Each camp has multiple states, and may not be available at all. If quests require single, specific camps in single, specific states, they will be very difficult to utilize and coverage will be poor. However, if quests can be adapted to different camps in different states, then the coverage and usage can be improved. We treat this as a constraint matching problem. A camp satisfies a quest if its state (see Table 1) unifies with the constraints placed on the quest. We consider two types of quest constraints and the challenges involved in relaxing them.

The first type of constraint is the objectives. These specify interactions with specific entities (mobs, nodes and/or NPCs) and locations (e.g. explore). Consider the very stereotypical example quest to kill bears in Figure 2.

We've had another bear attack! The animals in the forest to the East have been getting more and more aggressive. We need your help thinning their numbers.

Objective: Kill 10 Bears in the East forest.

Figure 2. Ubiquitous hunting quest.

This quest is constrained to a specific mob type and number, in a specific location. Variablizing those factors makes the quest more widely applicable. The names of entities and locations can be substituted in the discourse. Valls and Ontañón (2011) demonstrate techniques for repairing text following such substitutions. In addition, the number of kills required can be scaled to the available mobs. But objective constraints are not enough to guarantee that a camp is appropriate for a quest. The second type of constraint is the semantics of the discourse references. It makes perfect sense to substitute the bears with wolves, but not with bandits. Further, it makes sense to have bears in a forest, but not in a lake or a town. These are fairly general and straightforward relationships that can be captured by hierarchical type information and a small number of relevant, first order predicates. This type of information, though very difficult to generalize, is a minor addition to the cost of creating locations and entity types. Figure 3 shows example constraints in predicate calculus. With this information added, the example quest could be

constrained to Animal type mobs, and locations that have a known found-in fact for the chosen mob type.

(is-a Bear Animal)	(found-in Bear Forest)
(is-a Wolf Animal)	(found-in Wolf Forest)
(is-a Bandit Humanoid)	(found-in Bandit Forest)
...	...

Figure 3. Entity type and relationship information.

This system of entity substitution can be further expanded to other types of objectives. Consider the alternative quest in Figure 4, where killing the bears is only a means to an end. The item needed can be substituted as long as an entity known to produce the item is as well.

Winter's coming and we're short supplies. We could certainly use a hand. If you're willing to help, we need pelts from the bears nearby to make new winter blankets. You'll find them in the forest to the East.

Objective: Obtain 10 Pelts from Bears in the East forest.

Figure 4. Alternative hunting quest.

The quest discourse may also describe camp specifics that are not part of the objectives: a specific location feature, prop, camp size or conflict state. Like entity information, these constraints can be annotated on a quest to enforce narrative coherence during camp selection. Consider the examples in Table 2.

Sample Discourse Text	Constraints
"...attacked the mine and are stealing the gold!"	(holds ?location Gold)
"...polluting the water with their drilling machines..."	(has-prop ?camp Drill)
"...huge force of zombies has taken over..."	(size ?camp Large) (stage ?camp Established) (in-conflict ?camp None)
"...we haven't given up on our home, but they've pushed us into a corner..."	(in-conflict ?camp1 ?camp2) (stage ?camp1 Foothold) (stage ?camp2 Established)

Table 2. Sample discourse constraints.

If a quest-required entity is not found in a camp, it could fail to match, or it could add that entity to the camp. Whether that is appropriate depends on the discourse constraints and world information. For example, consider a quest to obtain a sacred idol by interacting with an altar node. If the quest was suitably unconstrained, a camp of witch doctors or a ruin full of wildlife could equally add the altar as needed to make the match.

Strategies for Expectation Failure

When a player arrives at a camp to pursue a quest, they expect it to be as described, and to be able to complete the objectives there. In this section we discuss strategies for

dealing with expectation failure, and their impact on this extended quest model.

To respond to the failures in the field, the quest model must be updated to support *interventions*. This is the ability to update objectives and discourse while the player is out working on a quest. Recent game designs already use automatic turn-in to minimize back-and-forth travel. For example, Red 5 Studio's science-fiction title *FireFall* uses a constant stream of radio messages to automatically turn-in and deliver quests as objectives are completed. As discussed prior, *GW2* simply updates progress and turn-in on the player UI. In this model, we go one step further to allow not only turn-in, but also alteration of objectives.

The first strategy when a camp fails to match the quest that sent the player there is to redirect the player to a suitable camp. The matching process detailed in the prior section can allow other camps to satisfy the quest. A second strategy is to declare the quest complete even though the player has not yet completed all the objectives. This only applies when the unexpected camp satisfies the narrative motivation behind the objectives. For example, a quest to remove a camp can be called successful if the camp is defeated while the player is still fighting. But a quest to exact revenge on a camp cannot be called successful if the camp is not there. A third strategy is to simply declare the quest a failure. The applicability conditions for this third strategy are approximately the inverse of the conditions for the second.

To support these strategies, the quest model must add the authorial burden of intervention text for each quest, explaining why things have changed the way they have. In many cases, it is likely that a pool of generic text could be used (e.g. "Looks like the bears aren't a problem anymore.", "No Bandits here, but you see tracks heading North."). However, the more detailed the story, the more likely specific intervention text will have to be provided.

Players are likely to find interventions tiresome if they happen too often. A fourth strategy addresses this by allowing the current quest to be swapped for a quest from a different story that matches the camp at hand. Clearly this cannot be done freely with any quest, because story coherence would be lost. Also, authoring transitions between all pairs of stories would be intractable. Instead, the model includes two techniques to enable swapping for some cases. The first is delayed commitment using *intervention breaks*. Quests are written that specify a location to go to, but have only vague commitments as to what will be found there. On arrival, the next quest begins with an intervention event that gives the specific objectives for that camp. That intervention break provides an opportunity to switch from the current quest to a different one that also has an intervention break in it. For example, consider the quest chain excerpts (a) and (b) in Figure 5, each telling its own story. Each of the first quests has a

vague initial objective that moves the player to a location, and an intervention break to provide new objectives and complete the quest. Excerpt (c) shows what the quest progress looks like if the second part of (b) is swapped in at the break in (a). The second technique required to make this work is providing context-free alternative text coming out of intervention breaks. The second part of (b) contains a back-reference to a prior poisoning incident in the quest chain ("The Goblins appear to be making the poison they used on those wolves from that Stinkweed."). The author must provide alternative text that reflects no prior context ("From the smell of that cauldron, what they're cooking up isn't dinner") in order for it to be swapped into (c). In Doran and Parberry's (2011) analysis of 750 quests across four MMORPGs, they found that the vast majority of quest actions involved travel to a new location. Writing those travel segments with intervention breaks and context-free alternatives would greatly increase flexibility in the system. However, there is no guarantee that a quest chain that gets swapped out will be able to swap back in later at that same quest in the chain. A challenging problem for adaptive narrative is reordering quests in a chain such that the story can be told in full from multiple entry points.

Discussion

This framework and the models of adaptive quest narrative are works in progress. Flexible constraint matching of camps to quests and strategies for expectation failure are proposed to expose the challenges of text adaptation and begin working towards a playable experience. A significant other challenge is how the system should choose to instantiate camps and hand out quests to maximize expected coverage, usage and coherence over all players. This requires some prediction of what players will choose to do, especially when given a new quest. In order to collect real data on the choices players make regarding quest progression, we have developed a lightweight MMORPG-inspired game implementing the standard world and quest models. We will be running controlled groups of players through various scenarios to collect that data. Future development will implement the extended models described in this paper.

A major omission here is the dimension of character. Many quest stories have character and relationship elements that progress through the arc, which are only loosely related to the objectives and world situations. For example, the first or last line of a quest description is often used to express the NPC's opinion of the player character. Representing those progressions as quest-less chains and inserting them into quests would not be difficult. The challenge is that the main quest text itself is often flavored

<p>(a)</p> <p>...</p> <p>Giver: Mr. Town</p> <p>There's something wrong with those bears you were hunting. They shouldn't have been out in those numbers so close to us.</p> <p>I'm worried, [%NAME]. Would you help us again by checking their caves further into the woods?</p> <p><i>Mr. Town wants you to investigate deeper into the East woods.</i></p> <hr/> <p>Intervention: <Dynamic attack by camp sentry></p> <p>That certainly wasn't a bear.</p> <p>You notice some Town scouts to the North. Perhaps they know what's going on.</p> <p>Completer: Mr. Scout</p> <p>Mr. Town sent you? Great, we could use the help. Those Bandits down there are just advance spies for something much bigger.</p> <p>...</p>	<p>(b) (discovered Goblin poisoning wolves)</p> <p>...</p> <p>Giver: Mr. Town</p> <p>I've just been hearing reports of highly agitated wolves starting to attack people outside the town. This is a deliberate attack!</p> <p>I need you to find out what else they're up to out there, and quickly!</p> <p><i>Mr. Town wants you to investigate the North woods for Goblin activity.</i></p> <hr/> <p>Intervention: <Arrive at brewing camp></p> <p>The Goblins appear to be making the poison they used on those wolves from that Stinkweed.</p> <p><i>Kill Goblin harvesters and knock over that brewing cauldron, then report back to Mr. Town.</i></p> <p>Completer: Mr. Town</p> <p>Goblins brewing up who knows what evil in our own backyard...</p> <p>We haven't heard the end of this, [%NAME].</p> <p>...</p>	<p>(c)</p> <p>...</p> <p>Giver: Mr. Town</p> <p>There's something wrong with those bears you were hunting. They shouldn't have been out in those numbers so close to us.</p> <p>I'm worried, [%NAME]. Would you help us again by checking their caves further into the woods?</p> <p><i>Mr. Town wants you to investigate deeper into the East woods.</i></p> <hr/> <p>Intervention: <Arrive at brewing camp></p> <p>From the smell of that cauldron, what they're cooking up isn't dinner.</p> <p><i>Kill Goblin harvesters and knock over that brewing cauldron, then report back to Mr. Town.</i></p> <p>Completer: Mr. Town</p> <p>Goblins brewing up who knows what evil in our own backyard...</p> <p>We haven't heard the end of this, [%NAME].</p> <p>...</p>
---	--	---

Figure 5. Example quest stories with intervention breaks and context-free alternative text.

with the NPC's attitude towards the player character, making incoherent combinations very probable. Work on automatically detecting character traits (Mairesse & Walker, 2006) may be a good starting point here.

References

- Doran, J., Parberry, I. (2011) A Prototype Quest Generator Based on a Structural Analysis of Quests from Four MMORPGs. In *Proceedings of the Second International Workshop on Procedural Content Generation in Games*. Bordeaux, France.
- Ducheneaut, N., Yee, N., Nickell, E., Moore, R.J. (2006) Alone Together? Exploring the Social Dynamics of Massively Multiplayer Online Games. In *Proceedings of the SIGCHI*.
- Fairclough, C., Cunningham, P. (2003) A Multiplayer Case Based Story Engine. Dublin, Trinity College Dublin, Department of Computer Science, TCD-CS-2003-43.
- Gervás, P., Díaz-agudo, B., Peinado, F., Gervás, R. (2005) Story Plot Generation based on CBR. *Journal of Knowledge Based Systems*.
- Li, B., Riedl, M. O. An offline planning approach to game plotline adaptation. *Proceedings of AIIDE 2010*.
- Mairesse, F. and M. Walker. Automatic Recognition of Personality in Conversation. *Proceedings of HLT-NAACL 2006*.
- Pita, J., Magerko, B., Scott, B.. True story: dynamically generated, contextually linked quests in persistent systems. In *Proceedings of Future Play 2007*.
- Riedl, M., Li, B., Ai, H., Ram, A. Robust and Authorable Multiplayer Storytelling Experiences. *Proceedings of AIIDE 2011*.
- Riedl, M., Sugandh, N. (2008) Story Planning with Vignettes: Toward Overcoming the Content Production Bottleneck. In *Proceeding of ICIDS 2008*.
- Reeves, S., Benford, S., O'Malley, C., Fraser, M. (2005) Designing the spectator experience. In *Proceedings of CHI*.
- Sullivan, A., Grow, A., Chirrick, T., Stokols, M., Wardrip-Fruin, N., Mateas, M. (2011) Extending CRPGs as an Interactive Storytelling Form. In *Proceedings of the ICIDS 2011*.
- Valls, J. and Ontañón, S. (2011) Text Modification Methods for Natural Language Generation. Master's thesis. Universitat Autònoma de Barcelona.
- Winegarden, J., and Young, R. 2006. Distributed interactive narrative planning. In *Proceedings of the 2006 AAAI Spring Symposium on Distributed Plan and Schedule Management*.
- Young, R. M.; Riedl, M. O.; Branly, M.; Jhala, A.; Martin, R.; and Saretto, C. 2004. An architecture for integrating plan-based behavior generation with interactive game environments. *Journal of Game Development*. 1:51-70.